Forecasting Research

Forecasting Research Division
Technical Report No. 232

An Archive of Aircraft Meteorological Observations and Model Forecast Meteorological Parameters

by

G J Rickard

January 1998

Meteorological Office
London Road
Bracknell
Berkshire
RG12 2SZ
United Kingdom

Forecasting Research
Technical Report No. 232

An Archive of Aircraft Meteorological Observations and Model Forecast Meteorological Parameters

by

G J Rickard

January 1998

NWP Division
Room 344
Meteorological Office
London Road
Bracknell
Berkshire
RG12 2SZ
United Kingdom

© Crown Copyright 1998

Permission to quote from this paper should be obtained from the above Meteorological Office division.

Please notify us if you change your address or no longer wish to receive these publications.

An Archive of Aircraft Meteorological Observations and Model Forecast Meteorological Parameters

G.J. Rickard

Aviation, FP(OS), Met. Office, U.K.

January, 1998

ABSTRACT

An archive combining commercial aircraft reports and forecast fields from the global forecast model at the Met. Office is produced routinely. This report details the archive contents, and outlines present applications of the archive, namely for verification of model forecast winds and temperatures against aircraft observations, and for the generation of parameters relevant to the prediction of clear air turbulence.

INTRODUCTION

The Met. Office routinely stores both meteorological data measurements, and forecast meteorological parameters from model runs. The measurements of interest here come from aircraft observations which are stored in what was the Synoptic Data Base (SDB), and what is now the Meteorological Data Base (MDB). The forecast parameters are stored in so-called Fields Files. Anyone seeking to compare the two sets of data is therefore compelled to extract the information from each data base, and then marry them together.

The process of marrying the data is feasible. However, there is an operational short cut, and that is to access the forecast parameters from the model runs *before* they become archived themselves as Fields Files; this is a less expensive exercise than marrying all the information at a later date. And indeed, this is the method used here. The operational forecasts are accessed, and "relevant" copies of the forecast fields are made. At a "suitable" later time, all the aircraft observations for a given forecast period are extracted, married to the forecast field copies, the final output forming the archive.

The question then is; what is actually to be stored in the archive? Clearly it is impractical to store "everything", and also unnecessary since "everything" is stored anyway! The answer is to store "enough" so that at a later date there is sufficient information in order to meaningfully investigate the forecast meteorological situation in the vicinity of the aircraft observation. The amount of "enough" should also be such so as to make the savings in computing time worthwhile.

For the simple purposes of point by point verification of the fundamental prognostic variables, the cube of forecast grid points surrounding the aircraft observation would suffice. However, an important diagnostic pertinent to aircraft from the model is that of the presence of Clear Air Turbulence (CAT). The parameters relevant to the prediction of CAT require spatial derivatives both horizontally and vertically, with the horizontal derivatives performed by finite differences, and those in the vertical by fitting cubic splines. Taking the CAT into account therefore leads to the requirement for a "tower" of data in which the aircraft observation is embedded. The structure needs to contain enough points in the horizontal and vertical in order to produce accurate estimates of the derivatives.

The archive thus consists of a series of towers of data about each aircraft observation, each appended to the next, and ordered in time. Diagnostic programs can then be used to access and post-process the data in the archive as required.

1. The Present Archive

The present version of the forecast model works on a regular longitude-latitude grid in the horizontal, and uses a hybrid coordinate η in the vertical, where η is pressure-based. Historically, however, many of the archived variables from the model forecasts are stored on standard pressure levels; the parameter values on these pressure levels are obtained by vertical interpolation from the η levels, where the interpolation is linear in the logarithm of the pressure. The present archive is therefore on ten standard pressure levels in the vertical, comprising the pressures 950, 850, 700, 500, 400, 300, 250, 200, 150, and 100 hPa (mb). Most of the commercial aviation reports of interest lie in the range 300 – 200 hPa (at least in the cruise phase of a flight); the stored pressure levels therefore have at least two extra pressure levels either side of this range of interest which, given that cubic splines are fitted in the vertical, results in good estimates of the vertical derivatives in the 300 – 200 hPa range. This can be seen by comparing the output for vertical derivatives from the splines using greater numbers of pressure levels against those for the ten standard levels here. It is clear that the spline values for the first and second derivatives deviate most strongly for the outermost two points at either end of the range.

In the horizontal, sufficient points either side of the aircraft observation on the regular longitude-latitude grid are stored so as to allow the usual second order finite difference approximations to first and second derivatives feasible; the chosen solution consists of recording the 16 points from the forecast model surrounding the aircraft observation, comprising a 4×4 set of points from the longitude-latitude grid.

Given the spatial layout of the archive, the question remains as to its temporal structure. The operational nature of the forecast model means that only forecasts at specific times are accessible (presently at regular six-hourly intervals). The aircraft observations are not so neatly ordered! The question then arises as to how to meaningfully construct the meteorological situation at any given time from the six-hourly spaced forecasts. In the light of any better suggestion, the method used is simple linear interpolation in time.

The global model from which the forecast fields are obtained is operationally run twice a day, with the initial conditions (or analysis fields) at times $00Z$ and $12Z$ on the day of interest. Clearly, to be of use to the airlines, the forecast data provided must arrive before the flight planning process is complete. To that end, the forecast fields at $T+12$, $T+18$ and $T+24$ relative to a given analysis time are of most relevance; the forecast time is as short as is practicable so that the airlines can perform their processing of the data ahead of their schedule on forecast fields as close to the analysis fields as possible. Therefore, in order to cover a complete day, the archiving process has to run twice a day, once on the $00Z$ run and once on the $12Z$ run. For comparison, forecast fields at $T+0$, $T+6$ and $T+12$ are also recorded, so that two archives are routinely stored.

Finally, there is the question of which meteorological parameters to store. The choice is mainly determined by the parameters sufficient to provide a measure of CAT and other turbulence indicators. At present these consist of; the temperature, the horizontal winds u and v, each of which is stored on all the pressure levels; the vertical velocity w stored on levels 950, 850, 700, 500, 300, and 250 hPa; the CAT probability (determined from the empirical index deduced by Dutton (1980)), which is stored on levels 300 to 200 hPa inclusive; and a series of parameters that are specific to the *whole* column of air, namely the mean sea level pressure, the tropopause pressure, temperature, and height, the maximum wind level pressure, cloud amounts for high, medium, low, and convective cloud, the convective cloud base and top, and the maximum wind components for u and v.

2. Spatial Layout of the Present Archive

The forecast model that is presently accessed uses the so-called Arakawa B grid in the horizontal. This means that not all the prognostic variables are solved for at the same set of grid points, i.e., the intersection points of the regular longitude-latitude grid. In particular the velocity variables are solved for on a grid that is "staggered" with respect to the "full" grid on which the remaining variables are held. Figure 1 shows the physical layout of what we shall refer to as the staggered or U-grid (the dashed lines), and the

full or P-grid (the solid lines). Clearly, the P-grid spans the entire physical domain from 0 to 360 degrees in longitude, and from -90 to $+90$ degrees in latitude. The U-grid is displaced from the P-grid by half the respective grid spacings $d\lambda$ and $d\phi$ in both longitude and latitude. This means that the U-grid spans the entire space longitudinally, but only covers the globe latitudinally from $(-90 + d\phi/2)$ to $(+90 - d\phi/2)$.

In storing the output variables from the forecast model, a choice to label them in the horizontal has to be made so that the stored location n, say, of a variable can then be translated into a physical location in the horizontal space. On the P-grid the counting is such that location number one occurs at the location $(0, 90)$ degrees. The count then increases along each row (or constant latitude) until the end of the row; the next location starts at 0 degrees again, but on the row displaced $d\phi$ Southwards. The process is repeated until all points have been counted. The same system applies to the U-grid, except that grid point number one lies at $(d\lambda/2, +90 - d\phi/2)$. Hence, in terms of a two-dimensional labelling of each grid point by (i, j), say, i increases with *increasing* longitude, whereas j increases with *decreasing* latitude.

For the U-grid plotted in Figure 1, all 16 grid points shown are to be counted and stored. However, for the P-grid, the set of points along $\lambda = 0$ are clearly the same ones along $\lambda = 360$ due to periodicity. The values along $\lambda = 360$ are therefore not stored, so that the total stored is actually 20 in this example. In longitude, then, both the P-grid and the U-grid will have the same number of grid points, whereas in latitude the U-grid will have one less than the P-grid.

In describing the archive in section 1, it was noted that a 4×4 set of points surrounding the aircraft report are stored. The question is how these points are defined. The definition used is that the point $(2, 2)$ from the 4×4 set of points is that grid point directly above in latitude and to the left in longitude of the aircraft observation. This applies to *both* the P-grid and U-grid. Given the point $(2, 2)$ in terms of its actual (i, j) value on the global grid (and hence its *single* value n), it is then a matter of counting to extract the relevant 4×4 grid points. The two-dimensional storage locations (i, j) in the archive follow the same convention as for the global grids, with i/j increasing for increasing/decreasing longitude/latitude.

The presence of the U-grid means care has to be exercised in the relationship between the label of a grid point and its physical location in the horizontal plane. This is emphasised in Figure 2, where a section of the global grid surrounding an aircraft observation labelled by '×' is shown. Again the U-grid is shown with dashed lines, the P-grid with solid lines. Four

particular aircraft observations ob1-ob4 are shown, each within a quadrant defined by the U-grid and the surrounding solid lines of a single grid cell on the P-grid. Clearly, regardless of the choice of ob1-ob4, the physical location of the point $(2,2)$ on the P-grid remains the same. However, the converse is true for the U-grid; for each aircraft observation the *physical location* of the point $(2,2)$ on the U-grid is different. Using the U-grid labelling in Figure 2 we find that $(2,2)$ for ob1 is at (BI, BJ), and then at (CI, BJ), (BI, CJ), and (CI, CJ) for ob2, ob3, and ob4, respectively.

In terms of the present variables stored in the archive, those on the P-grid are the temperature, mean sea level pressure, the tropopause pressure, temperature, and height, cloud amounts for high, medium, low, and convective cloud, and the convective cloud base and top. The remaining variables are all on the U-grid.

3. Post-Processing of the Present Archive

As we have seen, the main archive consists of a concatenated list of towers of data, with each tower surrounding a single aircraft observation and comprising forecast data that is linearly interpolated in time from the forecast records that span the actual time of the aircraft observation. Each tower comprises data on pre-determined pressure levels in the vertical, and on 16 grid points in the horizontal, with physical locations dependent on the position of the aircraft observation relative to both the P-grid and U-grid. The aim of the archive is to allow study of the relationship between the observed and forecast meteorological variables. In this section two particular uses of the archive will be described.

3a. Wind and Temperature Diagnostics

For commercial aircraft, major factors affecting operational safety and efficiency are the in situ winds and temperatures. In particular, the impact of the winds on forecast flight times is an integral part of the overall operational cycle of the airlines. To that end, it is important to be able to build up a picture of the accuracy of the forecast winds and temperatures along aircraft flight paths. This can be neatly extracted from our archive. A post-processing program simply accesses the archive, strips out the wind and temperature information for each aircraft record, and stores the data in a suitable form for further post-processing.

A particular requirement for the wind and temperature diagnostics is to obtain the relevant values at the actual aircraft location. The post-processing routine must therefore spatially interpolate from the grid values in the archive to the aircraft position; this is

performed by bilinear interpolation in longitude and latitude in the horizontal plane, and by linear interpolation in the logarithm of the pressure in the vertical. In performing the bilinear interpolation, care must be exercised over the particular grid (P or U) that the variable is stored on; this clearly impacts on the value obtained after interpolation.

3b. CAT Diagnostics

The second major use of the archive is to perform studies of CAT indicators, i.e., the ability of parameterised forms of the major prognostic variables to predict the occurrence or otherwise of CAT for commercial aviation. The nature of the archive allows not only continued verification of the present indicator (the CAT probability) against the aircraft observations of turbulence, but also the chance to evaluate other indicators.

3b1. Turbulence Indicators

The design of any post-processing routine for this task clearly depends on the form of the indicators of interest. For the present work they are;

$$VORABS = \frac{\partial v}{\partial x} - \frac{\partial u}{\partial y} + f, \tag{1}$$

where the Coriolis parameter $f = 2\Omega\sin(\phi)$, where Ω is the average angular velocity of the Earth ($7.29 \times 10^{-5} radians\ s^{-1}$), and ϕ is the latitude of the aircraft observation,

$$DST = \frac{\partial u}{\partial x} - \frac{\partial v}{\partial y}, \tag{2}$$

$$DSH = \frac{\partial u}{\partial y} + \frac{\partial v}{\partial x}, \tag{3}$$

and

$$CVG = -(\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y}). \tag{4}$$

The parameters defined in equations (1)-(4) depend entirely on horizontal derivatives of the winds u and v. These can all be determined by suitable centred difference formulations. An indicator based on the above is that due to Brown (1973) in the form,

$$BROWN = \sqrt{(0.3) \times VORABS^2 + DST^2 + DSH^2}. \tag{5}$$

The definitions for the vertical derivatives depend on whether a strictly differential formulation is used, or whether differences across a layer of atmosphere are to be considered.

In particular, if we define $V = (u^2 + v^2)^{1/2}$ and $\mathbf{V} = (u, v)$, then we can write,

$$\left|\frac{\partial \mathbf{V}}{\partial z}\right| = \frac{\partial V}{\partial p}\left|\frac{\partial p}{\partial z}\right|, \tag{6}$$

where

$$\frac{\partial V}{\partial p} = \sqrt{\left(\frac{\partial u}{\partial p}\right)^2 + \left(\frac{\partial v}{\partial p}\right)^2}\,. \tag{7}$$

In this (strictly differential) form, the value of $|\partial \mathbf{V}/\partial z|$ can be determined at a given vertical point by obtaining the vertical derivatives $\partial u/\partial p$ and $\partial v/\partial p$ from cubic splines, and $\partial p/\partial z$ from the hydrostatic relationship. However, in some formulations of turbulence indicators, such vertical derivatives are expressed in terms of changes of a variable across a specific layer of thickness Δz, say, over which there is a pressure change Δp, with resultant changes of u and v of Δu and Δv, respectively. In this formulation, the equivalent to equation (6) would be,

$$VWS = \frac{\sqrt{\Delta u^2 + \Delta v^2}}{\Delta z}\,. \tag{8}$$

With sufficient vertical resolution in the forecast models, the difference between these formulations would be expected to be small. The values found by the differential form at the centre of a layer should differ little from those found by the jump in the values across that layer, especially with the interpretation that a value at the centre of a layer "actually represents" an average value that applies across the whole layer. However, the relative coarseness of even the most up-to-date models means care has to exercised with these definitions.

In particular, two turbulence indicators promoted by Ellrod and Knapp (1992) are,

$$TI1 = DEF \times VWS, \tag{9}$$

and

$$TI2 = (DEF + CVG) \times VWS, \tag{10}$$

where $DEF = \sqrt{DST^2 + DSH^2}$. Furthermore, the work by Roach (1970a, 1970b) on the so-called "Richardson number tendency" proposes that the analytic expressions be diagnosed in terms of the jumps across the layers. The Richardson number is expressed as,

$$Ri = \frac{RT(\Delta \ln P)(\Delta \ln \theta)}{(DV)^2}, \tag{11}$$

where

$$DV = \sqrt{\Delta u^2 + \Delta v^2}, \tag{12}$$

and θ is the potential temperature, here taken to be $\theta = T(100000/p)^{\kappa}$, where $\kappa = R/c_p$, R is the gas constant, c_p is the specific heat capacity of dry air at constant pressure, and the pressure p is measured in Pa.

The tendency can be written,

$$\Phi = RK1 + RK2 + RK3. \tag{13}$$

The terms comprising Φ are,

$$RK1 = \left(2 - \frac{1}{Ri}\right)\left(\frac{R|\nabla_p T|(\Delta \ln P)\cos\beta}{DV}\right), \tag{14}$$

where $|\nabla_p T|$ is the magnitude of the *horizontal* gradient of the temperature at constant pressure,

$$RK2 = DST\cos 2\alpha, \tag{15}$$

$$RK3 = -DSH\sin 2\alpha, \tag{16}$$

where α is the clockwise angle between $\partial \mathbf{V}/\partial p$ and North, and β is the angle between $\partial \mathbf{V}/\partial p$ and ∇T.

In contrast, the indicator proposed by Dutton (1980) is based upon an empirical index E evaluated using the strictly differential forms, where

$$E = 1.25 S_H + 0.25(S_V)^2 + 10.5, \tag{17}$$

where S_V is the vertical shear given by equation (6) in units of $ms^{-1}\,per\,km$, and the horizontal wind shear is $S_H = -(\mathbf{V} \times \nabla V^2)/(2V^2)$ in units of $ms^{-1}\,per\,100\,km$. Expanded out, the latter term can be written,

$$S_H = \frac{1}{V^2}\left(uv\frac{\partial u}{\partial x} - u^2\frac{\partial u}{\partial y} + v^2\frac{\partial v}{\partial x} - uv\frac{\partial v}{\partial y}\right), \tag{18}$$

which changes sign when evaluated in the Southern hemisphere. The value of E is then used to determine the CAT probability that is output by the forecast model.

3b2. Turbulence Indicators and Aircraft Reports

Given all the above forms for the turbulence indicators, the question then is how and where should they be evaluated such that the indicator produced by the forecast model on a

discrete grid can be related in a meaningful way to the aircraft observations, particularly since the latter are not constrained to report at the discrete grid locations. The most sensible approach seems to be to consider the value at each grid point to be representative of the volume of air surrounding that point, with the volume having dimensions of $\delta\lambda(\phi) \times \delta\phi \times \delta z$ with its centre at the grid point. Here $\delta\lambda(\phi) = R_E \delta\lambda_o \cos\phi$ is the longitudinal grid spacing at latitude ϕ, where R_E is the average radius of the Earth, and $\delta\lambda_o$ is the longitudinal grid spacing at the equator; $\delta\phi$ is the (constant) grid spacing in latitude; and δz is the vertical width associated with the layer of atmosphere in question.

The present archive is stored on the ten pressure levels noted earlier. The vertical width could therefore be simply defined by placing the boundaries of a layer at the average pressure value between each of the stored pressure levels. Since the aircraft report gives the local value of pressure, it is a simple matter to determine which volume of air the aircraft is within. On the other hand, vertical linear interpolation is performed using the logarithm of the pressure equivalent to an effective height. It is therefore also plausible to define the layer boundaries in terms of the average height between each of the stored levels. The aircraft report is converted to a height, and the assignment to an air volume follows. The methods are compared in Figure 3. The stored pressure and the equivalent heights (in feet) at each vertical grid point are shown in the top two frames. The bottom two frames show $Pdiff$ and $Hdiff$ at the layer boundaries, where $Pdiff$ is the pressure difference between the pressure averaged and that obtained from averaging the height and converting back to pressure, and $Hdiff$ is the height difference between the height averaged and that obtained from averaging the pressure and converting to a height. Clearly, both methods result in relatively small differences.

We therefore opt to define the layer boundaries on the basis of the average pressure, since the pressure is the natural coordinate in the vertical. The particular layout for the ten standard levels presently used is shown in Figure 4. The bold horizontal lines refer to the standard levels (labelled in hPa, with the number in brackets just a simple counter), and the dashed lines to the intermediate pressure levels. The layers of major interest centred on 300, 250, and 200 hPa are indicated by the diagonal lines. Any errors that may result from this choice will cause a particular aircraft report to be associated with the "incorrect" volume of air and hence the particular value of a turbulence indicator at the grid point at the centre of the air volume. If the resultant predictability of the indicators turn out to be markedly affected by this arbitrary choice, it may suggest a lack of robustness in the CAT diagnostics exists.

The next task is to decide whether the indicators should be found on the P-grid or the U-grid. Examination of the indicators chosen for our study reveals that they depend entirely on the winds u and v, and the temperature T. The pressure p is known everywhere, and the potential temperature θ depends on p and T. Most of the indicators also depend entirely on the winds, so it seems appropriate to store the turbulence indicators on the U-grid. Horizontal averaging allows the temperature values determined by the model on the P-grid to be transferred to the U-grid; thereafter the calculations can proceed entirely on the U-grid.

While this choice is entirely appropriate for the majority of the indicators, it is clear that the Richardson number defined by equation (11), and its associated tendency in equations (13)-(16), contain information in different combinations from both grids. In particular, the tendency depends not only on the Richardson number, but also the angles α and β which define angles relating the vertical gradients in the horizontal velocity to horizontal gradients in the temperature. It is not at all obvious as to the sensitivity of these parameters to the averaging of the temperature onto the U-grid.

In summary, then, the turbulence indicators are to be determined on the U-grid in the horizontal, and at each standard pressure level in the vertical (although in practice we focus on the levels 300 to 200 hPa). Layer boundaries in the vertical are taken to lie at intermediate pressure levels (the dashed lines in Figure 4) which are at simple averages of the standard pressures. To determine parameters at the intermediate pressure levels, linear interpolation in the logarithm of the pressure is used; as a result, the so-called jumps across the layer for a value at the standard pressure level can be performed.

An aircraft report now needs to be associated with the volume of air, as previously defined. Although the definition is with respect to the volume of air, in reality it is simply a matter of finding the nearest grid point (NGP) to the aircraft report in the horizontal, and then to determine within which layer the report lies in the vertical. Figure 2 shows that the NGP on the U-grid is *always* (CI, CJ) in the archived data. The values obtained for CI and CJ will actually depend on the position of the aircraft report, as detailed in section 2. Following the guidelines in section 2, however, allows correct values for CI and CJ to be found. Having identified the relevant U-grid points, it is then a relatively simple matter to average the corresponding temperature values from the P-grid points onto the U-grid. Since the aircraft records the pressure, the layer in which it lies is easily obtained. Thereafter, consistent finite difference coding will allow final determination of

the turbulence indicator that can then be checked in relation to the aircraft observations.

4. Future Archives

Although the CAT probability provided by the forecast model is at standard pressure levels, the calculation to obtain those values actually takes place on the model η levels. It is generally argued that the relative spacing and location of the η levels mirrors that of the pressure levels at the level of commercial aviation, and therefore calculations on the pressure levels will be representative of those on the η levels. Such an assertion may indeed be true; however, many of the turbulence indicators are critically dependent on the vertical gradients, so it seems an unnecessary chance to take, particularly as the forecast data on the η levels can be made available. And indeed, the least tampering done with the model data in terms of interpolation etc, the more representative the post-processed results will be of those produced by the actual model forecast.

Clearly, then, any future versions of the archive should be stored on the hybrid η levels. Having made such a decision, then particular questions have to be answered before the final form of the new archive can be determined. The first concerns the number of levels actually required. Certainly the levels chosen should span the upper air space from 300 to 200 hPa. Furthermore, there should be sufficient levels outside this range in order to produce estimates of the vertical derivatives at the relevant levels consistent with those produced by the forecast model. The simplest solution is to of course store *all* the levels. Whether there will be sufficient storage is of course another question.

Secondly, what form do the formulae for the indicators take on η levels? η levels do not represent constant height or pressure surfaces in general, so that horizontal derivatives should in principle take this into account. However, in the present archive, the pressure surfaces are assumed to be (locally) constant height surfaces also, so that there is no distinction made between horizontal derivatives at constant pressure or constant height. The same approximations could be made on the η levels, and treat them locally for the purposes of horizontal derivatives as both constant pressure and constant height. Again, the relative sensitivity of the turbulence indicators to these approximations should be tested, at least cursorily.

In the light of the above questions and their answers, how many variables need now to be archived in order to correctly obtain the indicators in the hybrid system? Given the nature of the approximations regarding the horizontal derivatives, the only missing information is that of the pressure at a given point on an η level; this can be found knowing

the surface pressure distribution. Therefore, the addition of this variable would complete the formulation of a future archive. However, the caveats regarding the derivatives should be borne in mind.

Acknowledgements

This Technical Report's author arrived well after the inception of the archive. It is therefore pertinent to acknowledge those who have dirtied their hands in the creation and exploitation of the archive, in particular J. Human, A. McHaffie, N. Halsey, and R. Lunnon. A future protagonist appears to be C. Bysouth.

References

Brown, R. 1973: New indices to locate clear air turbulence, *Met. Mag.*, **102**, 347-361.

Dutton, M.J.O. 1980: Probability forecasts of clear air turbulence based on numerical model output, *Met. Mag.*, **109**, 293-310.

Ellrod, G.P. and Knapp, D.I. 1992: An objective clear-air turbulence forecasting technique: verification and operational use, *Wea. Forecasting*, **7**, 150-165.

Roach, W.T. 1970a: On the influence of synoptic development on the production of high-level turbulence, *Quart. J. Roy. Meteor. Soc.*, **96**, 413-429.

Roach, W.T. 1970b: Reply to "A note on the paper 'On the influence of synoptic development on the production of high-level turbulence' by W.T. Roach", *Quart. J. Roy. Meteor. Soc.*, **96**, 760.

Figure Captions

Figure 1. Schematic global grid layout in longitude (λ) and latitude (ϕ) in degrees. The bold lines show the P-grid, the dashed lines the U-grid.

Figure 2. Horizontal layout of archived data relative to the aircraft observation ('$\times$'). The P-grid is again in the bold, the U-grid staggered.

Figure 3. Vertical distribution of data levels. The top graph shows the standard pressure levels and the second the equivalent heights for those levels. The lower two graphs show the pressure and height differences, respectively, obtained by averaging in pressure then converting to height, or by averaging in height and converting to pressure (see text).

Figure 4. Labelling of vertical levels in the present archive. The dashed levels are at the intermediate pressure points, and the hashed regions illustrate atmospheric layers associated with a standard pressure level.

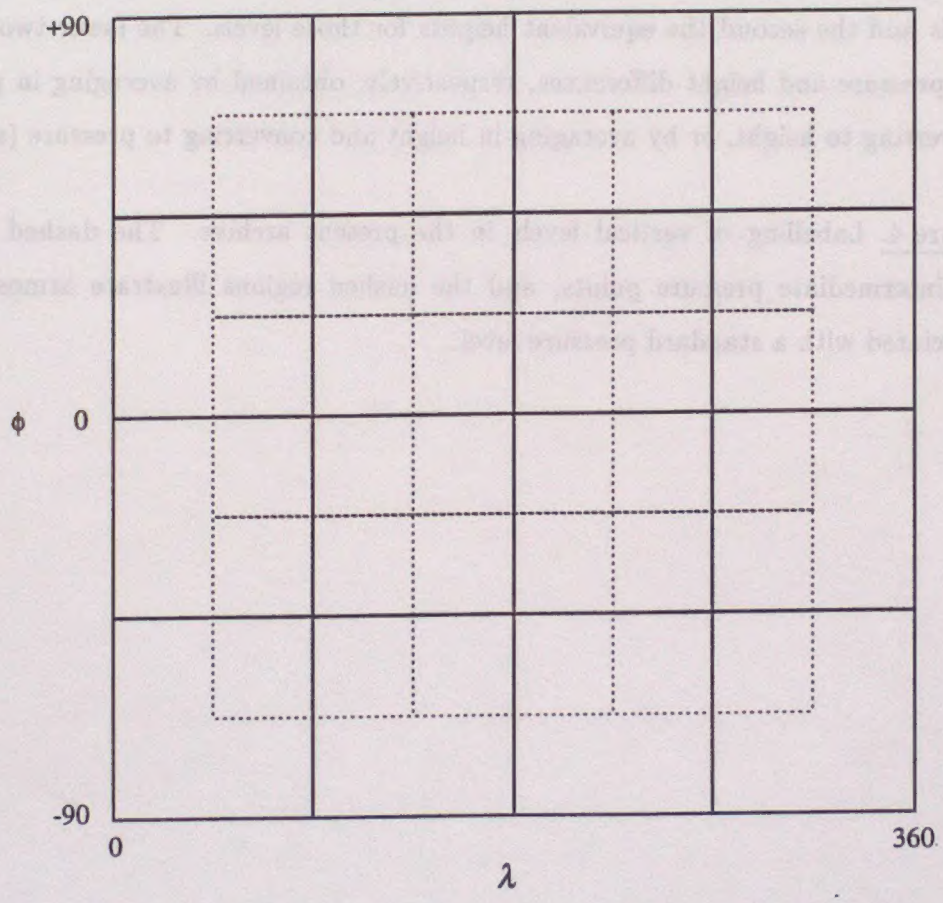

Fig. 1

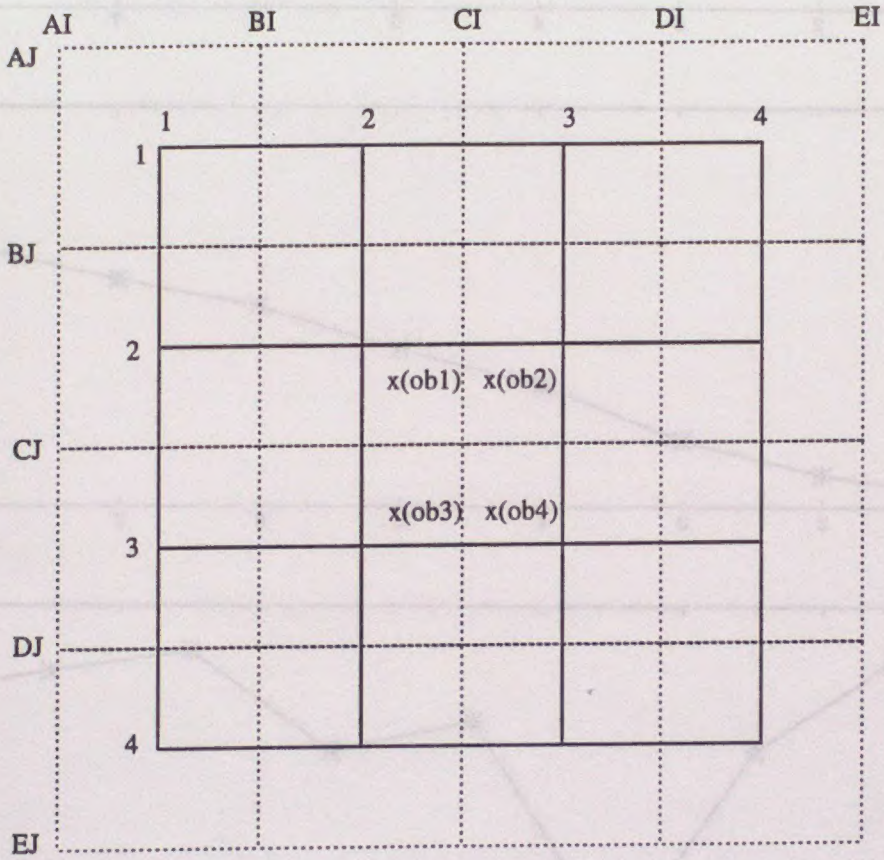

Fig. 2

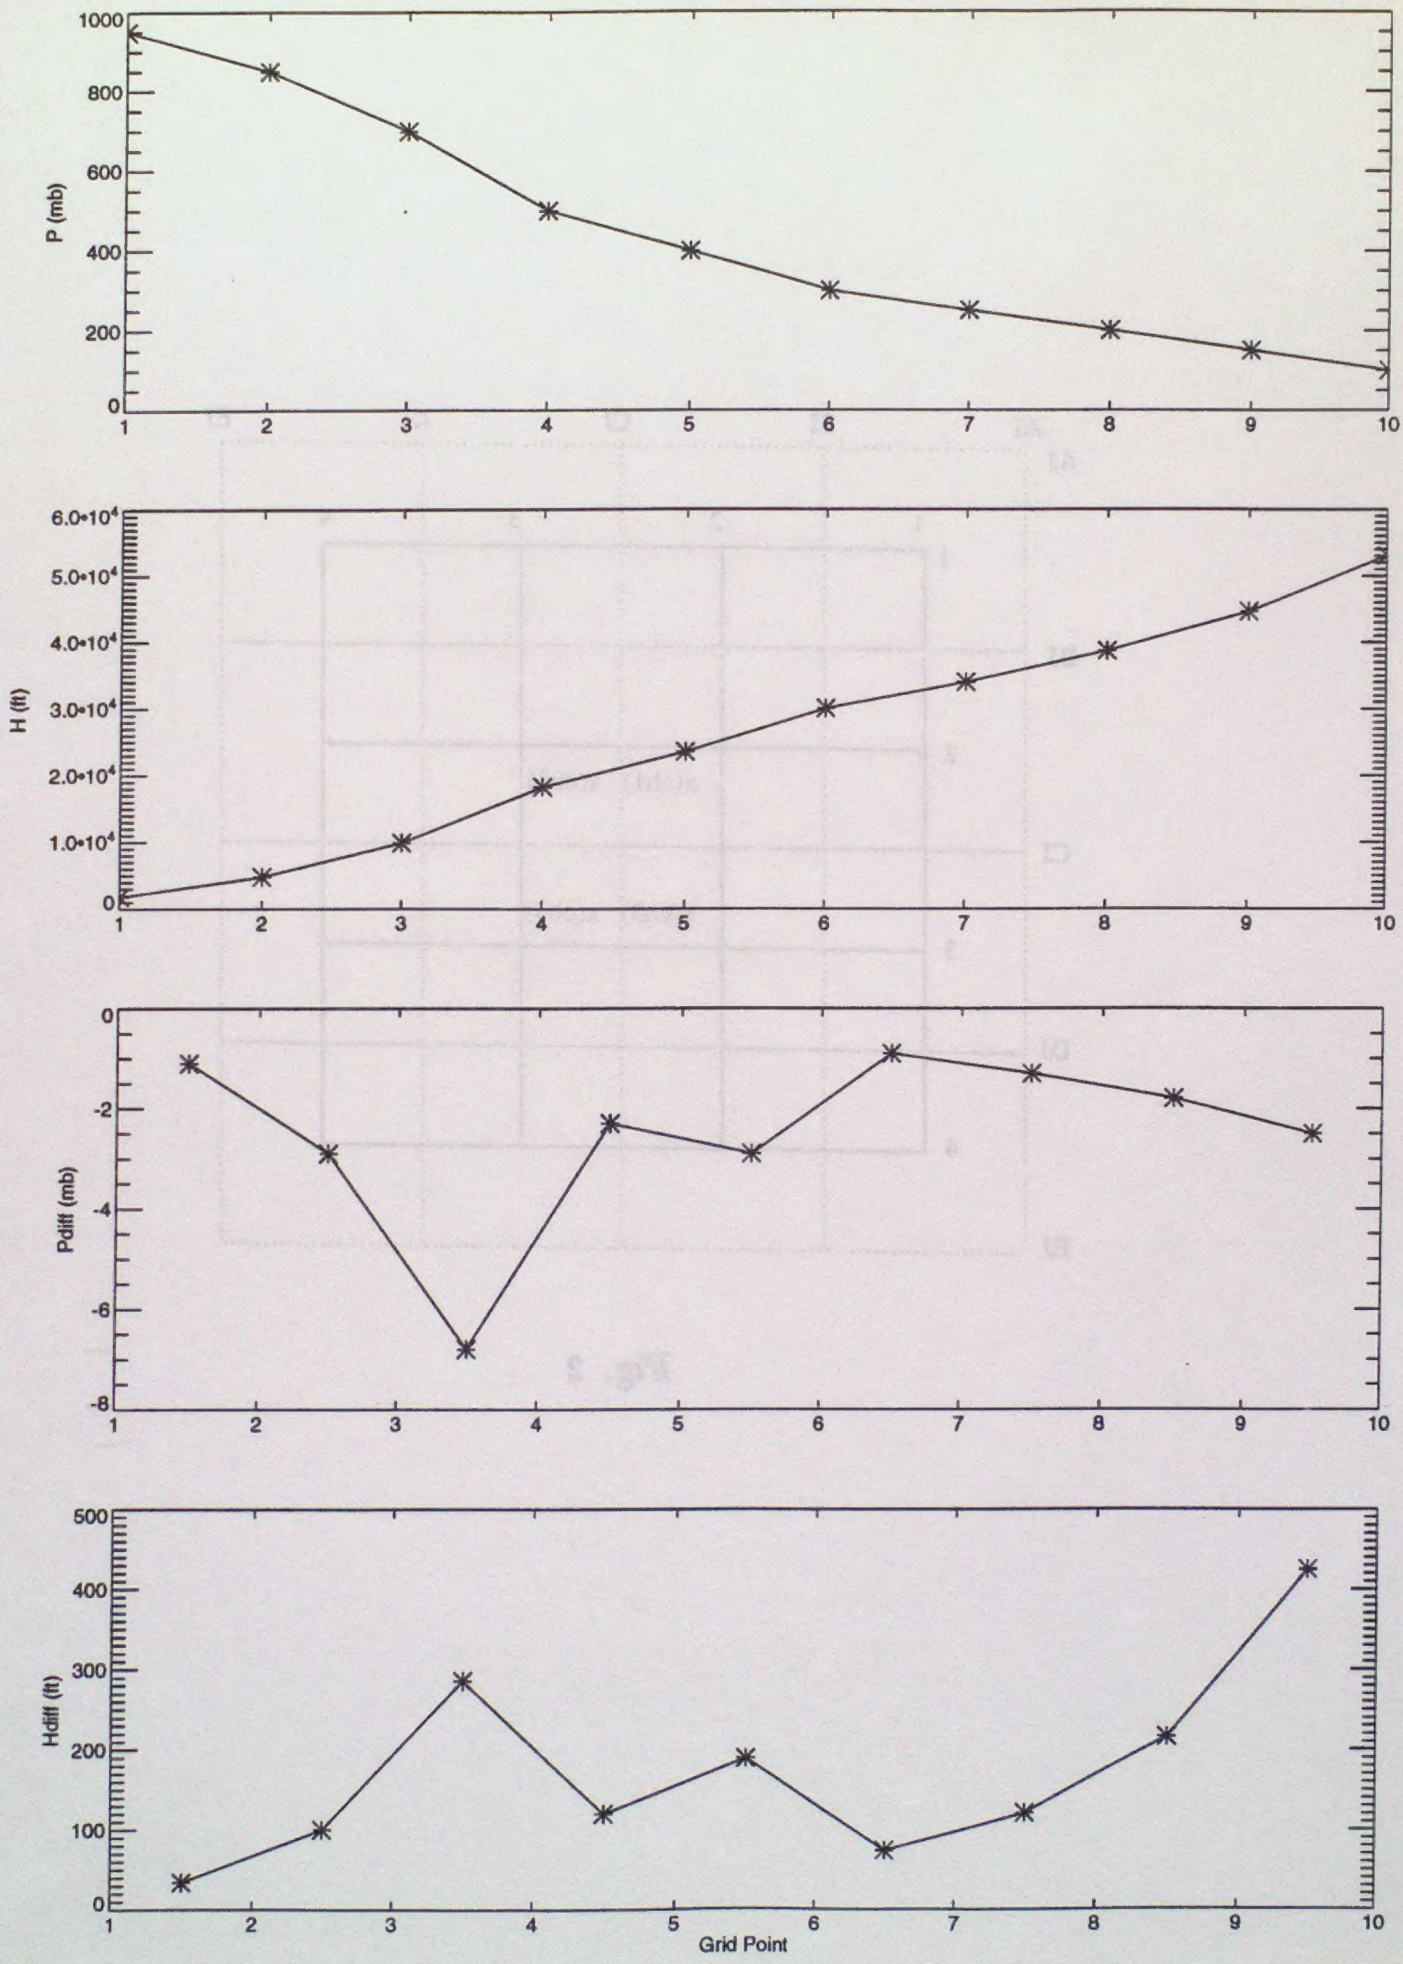
P (mb)
1000
800
600
400
200
0
H (ft)
6.0•10⁴
5.0•10⁴
4.0•10⁴
3.0•10⁴
2.0•10⁴
1.0•10⁴
0
Pdiff (mb)
0
-2
-4
-6
-8
Hdiff (ft)
500
400
300
200
100
0
1
2
3
4
5
6
7
8
9
10
Grid Point

Fig. 3

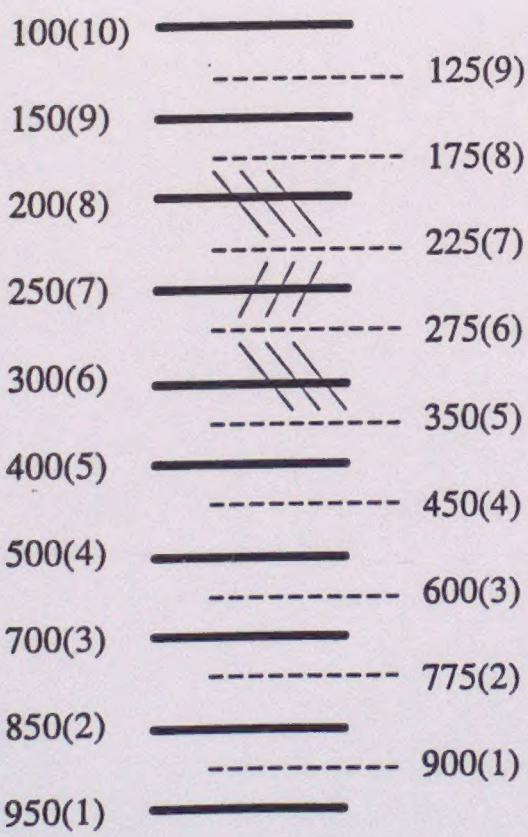
100(10)
125(9)
150(9)
175(8)
200(8)
225(7)
250(7)
275(6)
300(6)
350(5)
400(5)
450(4)
500(4)
600(3)
700(3)
775(2)
850(2)
900(1)
950(1)

Fig. 4